An Economic Analysis of Color-Blind Affirmative Action

Roland G. Fryer, Jr.*
Harvard University and NBER

Glenn C. Loury**
Brown University

Tolga Yuret***
Koc University

This article offers an economic analysis of color-blind alternatives to conventional affirmative action policies in higher education, focusing on efficiency issues. When the distribution of applicants' traits is fixed (i.e., in the short-run) color blindness leads colleges to shift weight from academic traits that predict performance to social traits that proxy for race. Using data on matriculates at several selective colleges and universities, we estimate that the short-run efficiency cost of "blind" relative to "sighted" affirmative action is comparable to the cost colleges would incur were they to ignore standardized test scores when deciding on admissions. We then build a model of applicant competition with endogenous effort in order to study long-run incentive effects. We show that, compared to the sighted alternative, color-blind affirmative action is inefficient because it flattens the function mapping effort into a probability of admission in the model's equilibrium.

> "Implementing race-neutral programs will help educational institutions minimize litigation risks they currently face. . . . If we are persistent in implementing race-neutral approaches, the end result will be to fulfill the great words of Dr. Martin Luther King Jr., who dreamed of the day that

*Department of Economics, Harvard University, and NBER. Email: rfryer@fas.harvard.edu.
**Department of Economics, Brown University. Email: glenn_loury@brown.edu.
***Department of Economics, Koc University. Email: tyuret@ku.edu.tr.

We are grateful to Ian Ayres, Richard Brooks, Avinash Dixit, Matthew Jackson, Lawrence Katz, Kevin Lang, Steven Levitt, Adrianna Lleras-Muney, Debraj Ray, and Andrei Shleifer for helpful comments and suggestions. We also thank seminar participants at numerous universities. Patricia Foo, Laura Kang, Eric Nielsen, and Emily Oster provided truly exceptional research assistance. Financial support is gratefully acknowledged from the Andrew W. Mellon Foundation and the Carnegie Corporation. Fryer thanks the Institute of Quantitative Social Science for generous support.

The Journal of Law, Economics, & Organization, Vol. 24, No. 2,
doi:10.1093/jleo/ewm053
Advance Access publication December 24, 2007

all children will be judged by the content of their character and not the color of their skin.''
—US Department of Education. *Race-Neutral Alternatives in Postsecondary Education: Innovative Approaches to Diversity*, Washington, D.C.: U.S. Department of Education, Office of Civil Rights March 2003, pp. 7, 40.

1. Introduction

The legal and political climate has shifted dramatically over the last decade on the issue of racial affirmative action. Accordingly, a number of institutions have begun to reformulate their policies—particularly in higher education. The states of Texas and Florida now guarantee admission to their public university systems for all in-state high school students graduating in the top 10% and 20%, respectively, of their senior classes.[1] In the wake of Proposition 209—a 1996 ballot initiative that banned racial affirmative action in California—public higher education officials there have substantially revised admissions practices.[2] Some private institutions have even decided to no longer require that applicants submit standardized test scores.[3] A number of scholars and policy analysts have urged elite colleges and universities to rely more on the socioeconomic background and other nonracial, nonacademic characteristics of prospective students when assessing their applications.[4]

Many justifications can be offered for these changes in admissions practice, but a primary factor would seem to be the desire to enhance racial diversity among the admitted without recourse to the use of explicit racial preferences. For this reason, we call these types of policies "color-blind affirmative action,"

1. In 1996, the state of Texas was ordered by a federal court to eliminate all race-conscious affirmative action in university admissions decisions [see Hopwood v. Texas, 78 F.3d 932 (5th Cir. 1996)]. The Texas legislature responded to Hopwood by passing House Bill 588, which guarantees Texas public high school students who graduate in the top 10% of their class admission to any Texas public college or university.

In February 2000, at the request of Governor Bush, the Florida State Board of Education banned consideration of race in admissions decisions for the state's higher education institutions. Florida's percentage plan, the Talented 20 program, took effect in August 2000. Under this plan, students who graduate from Florida's public high schools in the top 20% of their class, complete 19 specific academic credits, and take an scholastic aptitude test (SAT) or American College Test are guaranteed admission to one of 11 state universities, although not necessarily admission to the institution of the student's choice.

2. The so-called "Eligibility in the Local Context" policy was implemented in California in the fall of 2001. This program guarantees that the top 4% of each high school graduating class in the state will be admitted to one campus in the university system. For students admitted in the fall of 2002, the University of California system implemented a "comprehensive review" policy, which permits each campus to set admissions standards based on 10 academic and 4 nonacademic supplemental criteria, two of which may relate to socioeconomic status.

3. Mount Holyoke College, for example, has abrogated that requirement, although committing itself to admit some of the applicants who do not submit scores.

4. Richard Kahlenberg (1996) is perhaps the most prominent advocate of so-called class-based affirmative action policies.

in contrast to the more conventional, "color-sighted" affirmative action policies. Under color-sighted affirmative action, selectors give an explicit preference to individual applicants from some targeted racial group. A commitment to color blindness prohibits such behavior. Even so, group-preferential goals can still be pursued tacitly by exploiting knowledge of differences between the race-conditioned distributions of nonracial traits in the applicant population.[5]

In this article, we undertake a theoretical and empirical evaluation of the limits of "race-neutral approaches" such as those advocated by the US Department of Education in the passage quoted above. We are particularly concerned with the question of whether and the extent to which a widespread shift toward color-blind affirmative action might be expected to impair the efficiency of resource allocation in higher education. The answer to this question is of considerable importance for public policy.[6]

There are two distinct ways in which color-blind affirmative action is inherently inefficient. First, in the short-run, when the distribution of traits in the applicant pool may be taken as given, all affirmative action policies yield lower expected performance among the selected than does laissez-faire. This is due to the fact that, under laissez-faire (i.e., in the absence of any affirmative action policy), every admitted applicant is anticipated to perform better than any rejected applicant, which by definition cannot be true under any form of affirmative action. But, color-blind affirmative action is particularly inefficient in the short-run, in the sense that its performance is always dominated by the best color-sighted affirmative action policy calibrated to achieve the same group representation goal. This is so because the nonracial factors which best promote selection from a targeted group are necessarily different from the nonracial factors which best predict postselection academic performance—otherwise, some form of affirmative action would not be needed in the first place.[7]

Second, color-blind affirmative action is likely to be inefficient over the longer run as well, when one considers how the distribution of traits presented by

5. Obviously, introducing a purely random element to the selection process can also raise the yield from any group that is statistically underrepresented in the pool of admittees. This point is stressed by Chan and Eyster (2003). However, one important contribution of this article is to show that the options available to selectors for engaging in color-blind affirmative action are much broader that the simple use of randomization in the selection process.

6. It will be obvious in what follows that the ideas studied in this article are of quite general relevance. Color-blind affirmative action arises in many areas of public policy having nothing to do with enhancing racial diversity. For example, a powerful legislator may want to influence the formula specifying how some public benefit will be distributed among jurisdictions, with an eye toward benefiting his own constituency without appearing to be doing so. More generally, category-blind preferential policies can be used to pursue many group-redistributive goals (among population segments defined in terms of age, religious belief, gender, health status, region, nationality, and so forth), when decision makers wish to avoid the appearance of playing favorites. We elaborate on this point in the Conclusion.

7. The short-run efficiency of color-blind affirmative action depends solely on how well one can proxy for race with other observable characteristics—and how these characteristics relate to performance.

applicants will shift in response to the incentives created by colleges' admissions policies. Color-blind policies work by biasing the weights placed on nonracial traits in the admissions policy function so as to exploit the fact that some traits are relatively more likely to be found among the members of a preferred racial group. So, color-blind policies necessarily create a situation where the relative importance of traits for enhancing an applicant's prospects of being admitted diverges from the relative significance of those traits for enhancing an applicant's postadmissions performance. We show below that this is never the case under optimal color-sighted policy. Thus, to the extent that color-blind preferential policies distort applicants' decisions to acquire performance-enhancing traits prior to entering the selection competition, additional inefficiencies will emerge.

Our approach is simple and transparent. The central object of our analysis is what we call the "admissions policy function," which we imagine to be chosen by a college or university. Given the applicant pool, this function maps each applicant's "profile" into a probability of admitting that applicant. An applicant's profile is merely a list of that applicant's "score" along a number of dimensions, not all of which need be directly related to academic achievement.[8] An admissions policy function is said to be "color blind" if, other things being equal, the probability of admission that gets assigned to a profile does not depend upon an applicant's race. (Likewise, a policy function that makes use of race is said to be color-sighted.) Since color blindness is an additional constraint on the admissions process, given any target rate of admission from a "disadvantaged" minority group the best color-blind policy meeting that target must perform less well, from a college's point of view, than the best color-sighted policy. Using the College and Beyond database, we examine data from matriculates at seven elite colleges and universities to understand the magnitudes involved. The analysis proceeds as follows:

First, in our simulation exercises we allow colleges to make hypothetical admissions decisions based on a vector of academic and nonacademic applicant traits. This permits us to deduce how, in the short-run, a broad reliance on color-blind affirmative action might lower selection efficiency and alter the relative weights given to various factors in the college admissions process—grades versus test scores versus socioeconomic background, for instance. (We wish to stress that our analytical apparatus is flexible enough to encompass all the aforementioned color-blind practices—percentage plans, voluntary test score submission, increased relative weight on non–test score criteria, preferential admission based on socioeconomic status—as well as conventional affirmative action policies, within a unified framework.)

8. In general, an applicant's chance of admission can be made to depend upon a host of factors. Conventional academic variables—test performance, grades in high school, recommendation letters, interview results, and so on—can be supplemented with information about an applicant's social background, life experience, geographic region of origin, extracurricular interests, and the like. We think of the specific variables used in an admissions policy function and the weights given to them as being chosen by the college or university in order to meet its admissions objectives.

Second, in a theoretical model with endogenous applicant effort, we study the ways that color-blind policies diminish the equilibrium incentives that applicants have to acquire traits valued by selectors. This gives some sense of the possible longer run efficiency costs of such policies. In the model, students anticipate the colleges' policies prior to applying for admission and make a binary, costly effort decision that affects the distribution of their academic qualifications. In the unique equilibrium under color-blind affirmative action, as the colleges' representation target approaches population parity the fraction of students choosing high effort approaches zero. By contrast, under color-sighted affirmative action, a goal of population parity can be achieved in equilibrium without vitiating students' effort incentives. Our principle conclusion is that to rely solely on race-neutral approaches to achieve greater racial diversity in higher education would be to risk some possibly serious, and negative, unintended consequences. (Of course there may be other, nonefficiency-related, reasons to forego the use of race in college admissions. However, such considerations lie beyond the scope of this strictly economic analysis.[9])

Much has been written on the pros and cons of affirmative action, especially in the labor market.[10] However, until quite recently there had been little attention given in either the theoretical or empirical literatures to resource allocation inefficiencies due to affirmative action in higher education.[11] Two recent contributions warrant to be mentioned. Chan and Eyster (2003) have independently made one of the observations which we stress here—namely, that a ban on affirmative action could induce colleges to use inefficient, nonracially preferential means to pursue their racial diversity ends. They study a constrained-optimal admissions problem for a college that values both student quality and racial diversity, that can rank students based on a

9. For a more extended, critical discussion of these race-neutral approaches—in the context of a specific legal dispute over the constitutionality of racial affirmative action at public universities—see the amicus curiae brief filed with the US Supreme Court in the case Grutter v. Bollinger involving the University of Michigan Law School (Loury et al. 2003).

We realize, of course, that efficiency is not the only concern when assessing the desirability or the legality of alternative affirmative action policies. However, under the current Supreme Court's standards of legal scrutiny, a racial preference can be permitted if it constitutes a "narrowly tailored" means of furthering a "compelling state interest." Thus, once the goal of enhanced racial diversity in college admissions is acknowledged to be a compelling one, efficiency considerations become relevant to the legal determination of whether a given policy has been narrowly tailored to advance that purpose. A grossly inefficient policy, relative to some feasible alternative that achieves the same racial representation goal, is not a narrowly tailored one. See Ayers (1996) and the related discussion in Loury et al. (2003).

10. Coate and Loury (1993) develop a theoretical framework for analyzing the incentive effects of affirmative action in the labor market. The article by Holzer and Neumark (2000) is a comprehensive and insightful review of the theoretical and empirical literatures on affirmative action.

11. The article by Datcher Loury and Garman (1993) is an exception. That article argues empirically that racial preferences in college admissions may induce an inefficient assignment of minority students to institutions (differentiated by their degree of selectivity). However, the evidence on this question is mixed. Using different data, Kane (1998) finds no support for the hypothesis of a detrimental mismatch for minority students due to (color sighted) affirmative action in college admissions.

one-dimensional measure of student ability, but that is enjoined from using racial preferences. They show that in their model the second-best optimal admissions policy generally involves randomization. Chan and Eyster conclude, as do we, that a ban on color-sighted affirmative action could end up lowering the average quality of the college's admitted class. However, their analysis is not comprehensive: It fails to take into account the fact that colleges can use nonracial proxies, and not just randomization, as a way to enhance racial diversity under color blindness. Moreover, they cannot address long-run efficiency issues at all because, unlike in the present study, their theoretical model treats applicant characteristics as exogenous.

In another recent article, Epple et al. (2003) take note of the fact that a prohibition on explicit affirmative action can be expected to alter a college's use of nonracial information in the admissions process. Their numerically simulated model complements ours by focusing on the supply side of the higher education market. They introduce a framework where colleges differ in their attractiveness to applicants and compete with one another for the most desirable students. They are thus able to address the important question (which we here ignore) of how the distribution of students across a quality hierarchy of colleges would be affected by a ban on explicit racial preferences. However, they also take the distribution of applicant traits to be exogenous. Overall, their analysis focuses on a different set of issue than those explored below.

The structure of thisarticle is as follows. Section 2 describes our empirical model used to implement color-blind affirmative action and presents estimates of the efficiency losses involved in the short-run. Section 3 develops a model of incentive effects with endogenous traits to illustrate the long-run consequences of the widespread adoption of color-blind affirmative action. Section 4 concludes by discussing how the methods developed in this article might be applied to other policy issues. There is a technical appendix which contains all formal results stated in this article.

2. Color-Blind Affirmative Action in the Short-Run

To fix ideas, consider a concrete example of how color-blind affirmative action might work. Suppose initially that a college wants to admit a certain fraction of its applicants while maximizing the expected performance of those admitted. Let expected performance be a linear function of standardized test scores and of extracurricular activities in high school.[12] It is clear, then, that this college should adopt the policy of admitting only those applicants whose expected performance exceeds some threshold, where this threshold has the property that the fraction of applicants exceeding it just equals the fraction the college desires to admit. In effect, this means that the "weight" the college gives to activities relative to test scores in its admissions policy function should equal

12. We focus in the present example on two variables likely to enter any college's admissions policy function without intending to imply that these are the only variables of interest.

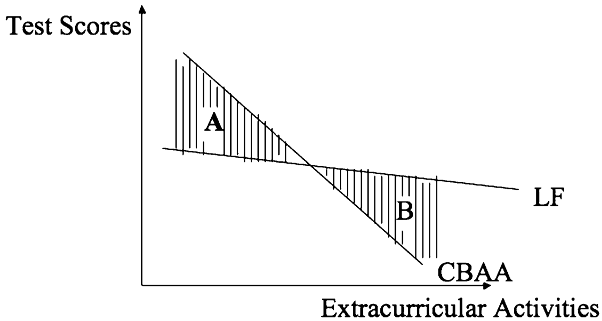

Figure 1. A Two-Attribute Example.

the ratio of the respective partial correlations of these variables with postadmissions performance. Now, suppose the college believes that following this threshold policy would lead to "too few" members of some racial group being admitted.

Imagine that the college wants to obtain a greater degree of racial diversity while continuing to be race blind in its treatment of individual applicants. Finally, suppose the college knows that among its applicants the distributions of activities within racial groups are much more similar to each other than are the corresponding distributions of test scores. Then the representation of the racial group with relatively lower (higher) test scores could be enhanced by setting the weight given to extracurricular activities relative to test scores in the admissions policy function above (below) the level warranted by the relative correlations of these variables with postadmissions performance. To introduce such a change in admissions policy would be to engage in the practice of color-blind affirmative action.

Figure 1 captures the intuition at work. The line segment LF represents a college's optimal admissions frontier under a policy of no affirmative action (call this laissez-faire). Applicants above the line are admitted with probability one, whereas those falling below the line are admitted with probability zero. The line segment CB represents the same college's admission frontier under a policy of color-blind affirmative action. The CB frontier is steeper than the LF frontier because we have supposed that extracurricular activities are more nearly equally distributed within racial groups than are test scores. The shaded area marked A in the figure depicts the set of applicants (with high test scores and low activities) who are rejected under CB, but who would have been admitted under LF. The area marked B shows the set of applicants (with high activities and low test scores) who are admitted under CB, but would have been rejected under laissez-faire. Because the college intends to fill a fixed number of seats, the number of students falling within each of these two areas is the same. Yet, because the conditional probability that an applicant belongs to the targeted racial group is greater given that the application falls in area B than it is given that the application falls in area A, this college can enhance racial diversity in a race-blind manner by raising the weight it gives to extracurricular activities relative to test scores when evaluating all applicants.[13]

2.1 An Empirical Model of Short-Run Affirmative Action

We now extend and formalize this example. Imagine that a college is to select an incoming class from a finite set of applicants. Let c denote the proportion of applicants to whom admission can be offered, $0 < c < 1$, and let r denote the target admissions rate for a disadvantaged minority group (relative to the size of the applicant pool). Let I be the set of all applicants, and take $i \in I$ to index a particular individual.

Suppose that each applicant belongs to one of two racial groups, and let $R_i \in \{1, 2\}$ denote the racial group membership of applicant i. Each application reports values for a bundle of nonracial traits (grades, social background factors, test scores, and the like). Let J denote the set of nonracial traits, with specific traits indexed by $j \in J$. Then, the ith student's application can be represented by the vector $(R_i; x_i)$, where $x_i \equiv (x_i^j)_{j \in J}$, and where x_i^j is the value which the application of student i reports for nonracial trait j. Moreover, the college's entire applicant pool can be represented by the (large) array $X = \{(R_i; x_i)_{i \in I}\}$.

Now, in general, an admission policy for the college in this setting associates with every applicant pool an array of probabilities specifying the chance that each applicant in the pool will be admitted. Let A_i be the probability of admitting applicant i, $0 \leq A_i \leq 1$. Then the college's admissions problem is to associate with each applicant pool, X, a vector of admission probabilities, $A(X) = (A_i)_{i \in I}$, so as to maximize the expected academic performance of the admitted class, subject to its capacity and racial representation constraints.

Let p_i be the college's expectation of the academic performance of applicant i. We assume that this expectation can be expressed as a linear function of the applicant's nonracial traits:

$$p_i \equiv [\text{Expected performance} \mid x_i] = \beta \cdot x_i = \sum_{j \in J} \beta_j x_i^j$$

for some vector of coefficients, β. In addition, to the extent that the nonracial traits are distributed differently within the racial groups, a college could use this fact to predict the racial group membership of any applicant presenting

13. Note that this enhanced racial diversity is achieved at the cost of admitting a lower performing class on average, since the expected performance of every applicant in B is lower than that of any applicant in A.

Furthermore, suppose a college were to make reporting test scores optional for its applicants, although committing itself to admitting a certain fraction of its incoming class from the set of students electing not to submit scores (as Mount Holyoke College has, in fact, recently done.) In light of the incentives thereby created for applicants to selectively report their test scores, this too would be a color-blind policy that, although not explicitly preferential to any racial group, could be expected to result in more students from low-scoring groups being admitted.

a particular vector of nonracial traits. Again, adopting a linear specification of this relationship, we assume that:

$$r_i \equiv \Pr[R_i = 2 \mid x_i] = \gamma \cdot x_i = \sum_{j \in J} \gamma_j x_i^j$$

for some vector of coefficients, γ.

In what follows, we take it that the vectors of coefficients, β and γ, are known to the college and enter as parameters in its calculation of an optimal admissions policy. We use our data on matriculates at several selective institutions to estimate these coefficients. We then use these estimates to simulate what optimal admissions policies might look like under various regimes at these colleges and to evaluate their performance.

We examine the implications of three distinct policy regimes: laissez-faire (LF), color-sighted affirmative action (CS), and color-blind affirmative action (CB). We begin by considering the following simple linear program:

$$\max_{[A_i]_{i \in I}} \left\{ \left(\frac{1}{c}\right) \sum_{i \in I} A_i p_i \right\}, \quad \text{subject to the following three constraints:}$$

$$\text{(i)}\ A_i \in [0,1],\ \ i \in I, \quad \text{(ii)}\ \frac{1}{|I|}\left\{\sum_{i \in I} A_i\right\} \leq c, \quad \text{(iii)}\ \frac{1}{|I|}\left\{\sum_{i \in I} A_i r_i\right\} \geq r.$$

The maximand above is the anticipated average performance of the admitted class. Constraint (i) restricts the A_i to being probabilities, (ii) is a capacity constraint, and (iii) is the affirmative action representation constraint.

A college's optimal CB admissions policy must solve this linear programing problem. An optimal LF policy solves the same problem, but without constraint (iii). An optimal CS policy can be derived by first partitioning the applicant pool by race and then solving parallel linear programs for each group, analogous to the LF version of the program above, but with the group-specific capacity constraints $r_2 = \frac{r}{1-\lambda}$ for group 2 and $r_1 = \frac{c-r}{\lambda}$ for group 1.

Solutions for the optimization problems implied by the three policy regimes are easily derived. Under the LF regime, one simply orders applicants by their expected performance, admitting the proportion c with the higher values of p_i. That is, for some number μ (the Lagrangian multiplier on constraint (ii) above), we have that:

$$A_i^* = \begin{cases} 1 & \text{if } \beta \cdot x_i > \mu \\ 0 & \text{if } \beta \cdot x_i < \mu. \end{cases}$$

Here μ must be chosen in such a way that constraint (ii) holds with equality.[14]

14. If $\beta \cdot x_i = \mu$, then A_i^* might need to lie strictly between zero and one for the capacity constraint to hold with equality.

Under the CS regime, there will be separate thresholds for the racial groups. So, for a pair of numbers μ_1 and μ_2, with $\mu_1 > \mu_2$, we have:

$$A_i^* = \begin{cases} 1 & \text{if } \beta \cdot x_i > \mu_{R_i} \\ 0 & \text{if } \beta \cdot x_i < \mu_{R_i}. \end{cases}$$

Here the μ_1 and μ_2 are to be chosen such that selection rates for the two groups are consistent with the capacity and representation constraints holding as equalities.

Under the CB regime, a Lagrangian multiplier on constraint (iii) alters the admissions policy relative to LF because nonracial traits are now to be valued both for their association with prospective academic performance and for their ability to predict an applicant's race. Thus, the optimal CB policy is characterized by two numbers θ and μ' such that:

$$A_i^* = \begin{cases} 1 & \text{if } [\beta + \theta\gamma] \cdot x_i > \mu' \\ 0 & \text{if } [\beta + \theta\gamma] \cdot x_i < \mu', \end{cases}$$

where μ' and θ are such that constraints (ii) and (iii) above hold as equalities.

This formalization captures nicely the ideas about color-blind policy mentioned in the introduction. Let j and k be two traits (e.g., extracurricular activities and test scores). Under LF and CS regimes, the college's marginal rate of substitution between traits j and k as reflected in the admissions policy function, denoted by $\mathrm{MRS}_{j,k}$, is equal to the relative importance of these traits in forecasting student performance:

$$\mathrm{MRS}_{j,k} = \frac{\beta_j}{\beta_k},$$

whereas, under the CB regime, the rate of substitution between traits j and k that holds constant the probability of being admitted is given by:

$$\mathrm{MRS}_{j,k} = \frac{\beta_j + \theta\gamma_j}{\beta_k + \theta\gamma_k}.$$

These substitution rates are the signals sent out to applicants about the relative value of various traits in the admissions process. To the extent that the magnitude (and even the sign!) of such substitution rates is altered when color-blind means are used to pursue color-conscious ends, the incentives applicants have to acquire the relevant traits might be badly misaligned. We will now use our data on student characteristics at selective public and private colleges and universities in the United States to examine how color-blind affirmative action might be expected to play out in practice.

2.2 Simulating the Short-Run Impact of Color-Blind Affirmative Action

To apply the foregoing analysis in the context of college admissions in the United States, we will use actual student profiles from the matriculating classes (entering college in 1989) of seven selective institutions (four liberal arts colleges, labeled "College A" through "College D" and three research universities, labeled "College E" through "College G" in what follows). We conduct hypothetical admissions experiments, supposing that the colleges in question would have had to "admit" only a fraction as many students as were, in fact, admitted.[15] Their imagined selection problem is to choose which students to retain, and which not, from among the actual matriculates. The affirmative action goal is to maintain the original proportion of minority students in this reduced class. We estimate the loss of efficiency in selection that results from the imposition of the requirement to be color blind in the selection process, given this racial representation goal. We also look at the nature of the constrained-optimal color-blind admissions policies that emerge.

Thus, the capacity constraint for all colleges in these empirical exercises, unless otherwise noted, is $c = 0.5$.[16] However, the affirmative action representation target under CS and CB varies from college to college, since the admissions policy maker seeks to maintain the same percentage of blacks among the selected students as had obtained in the original class and that percentage varies across college. Employing the framework just discussed, we model the constrained policy choices in each regime as linear optimization problems: an admissions policy is chosen, given the distribution of applicant traits and subject to capacity and representation constraints, so as to maximize the anticipated average academic quality of the admitted class. Once solutions for these linear programs are in hand for each college, we can compare the performance of the best admissions policy under each of the three regimes and take note of how the constrained-optimal color-blind policy attains its goal through an artful choice of racial proxies.

2.3 Data Description and Empirical Implementation

The College and Beyond database is remarkably rich—containing student-level administrative data on college performance as well as information on admissions and transcript records of 93,660 full-time students who entered 34 colleges and universities in the fall of 1951, 1976, and 1989 (see Bowen and Bok [1998] for a complete description). For the purposes of this article, we restrict our attention to students from seven institutions in 1989. Our selection criterion is based solely on the availability of relevant data. Section 4 (the data appendix) describes how we combined and recoded some of the College and Beyond variables we use in our analysis.

We employ three academic variables—scholastic aptitude test (SAT) math score, SAT verbal score, and high school rank—and six socioeconomic background variables—mothers education; fathers education; the median household

15. The ideal data for our thought experiment would come from a double-blind experiment in which neither a college's faculty nor its students would know what affirmative action treatment (LF, CS, or CB) they were under. This would not only solve the obvious selection problems but also alleviate any externalities that may arise from faculty or peers knowing that blacks were admitted under a particular regime. Unfortunately, however, these data do not exist.

16. We will test the robustness of our results to this assumption in the next sections.

income of each student's zip code; the percent black, Hispanic, and Asian in each students zip code; and whether or not each student is related to an alumnus (i.e., a legacy).[17] Only one SAT math score and one SAT verbal score was recorded for each student, even if the student took the test multiple times. Information is not available pertaining to which SAT score the institution reported. Parental education information was drawn from the student's college application. Questions involving parental education varied greatly from university to university. To account for this, we aggregated the data into two categories: college degree holder or not, independently, for mother's and father's education. Zip income was calculated by obtaining each student's residential zip code from College and Beyond data set and imputing the median household income by zip code from the 1990 Census. Percent racial mix in zip code was gleaned similarly. Legacy status was given if any family member was an alumnus of their university. If legacy status was unknown, we considered the student a nonlegacy.

We use linear regression analysis to associate these variables with the expected class rank after 4 years of matriculation achieved by the students in the sample (whose grade histories were available from the administrative records of the participating institutions). That is, we estimate models of the form:

$$\text{GPA\_RANK}_{ij} = X_i\beta_j + \varepsilon_{ij}, \tag{1}$$

where GPA_RANK_{ij} denotes the cumulative GPA of individual i at college j and X captures the various academic and social variables described above. We also use these covariates in a linear probability specification to estimate the conditional likelihood that a given student is of a given race, given his or her nonracial characteristics, of the following form.[18]

$$\text{Race}_{ij} = X_i\gamma_j + \upsilon_{ij}. \tag{2}$$

To simulate the short-run impact of color-bind affirmative action, we estimate equations (1) and (2) on the full sample of each college, individually.[19] With the estimated coefficients, we can conduct our hypothetical admissions experiments. Under LF, a college will admit a student only if $\sum_{j\in J}\beta_j x_j^i \geq \mu$. Under CS, a college will admit a student only if $\sum_{j\in J}\beta_j^R x_j^i \geq \mu_R$, $R \in \{1,2\}$. Under CB, a college admits only if $\sum_{j\in J}[\beta_j + \theta\gamma_j]x_j^i \geq \mu'$. The latter requires us to solve a simple linear program to obtain the weights, θ, given to covariates.

17. Unfortunately, at some colleges, family income information was only available for those students who applied for financial aid. So, we have used the median household income in the zip code of residence at the time of application as a proxy for socioeconomic status.

18. Note that, although the estimated parameters β and γ are intended to apply to a college's applicant pool, the data on which our estimates are based come from matriculates, not applicants, at the various colleges. As a consequence, the coefficient estimates presented here could be seriously affected by problems of selection bias.

19. One could also run the regressions within race. We run everything on the full sample so as to not confuse the inefficiency of race blindness that occurs because colleges shift emphasis away from academic traits and put more weight on social traits from the fact that they just have a better regression specification under laissez-faire and race-conscious admissions. That is, we use a specification on the full sample that is constant under all three regimes to isolate the desired effect.

2.4 Selection Problems

The ideal experiment to test the short-run inefficiency of alternative affirmative action policies would randomly distribute such policies across a wide swath of universities. By comparing the pre- and postquality of the admitted classes across similar universities, one can get an unbiased estimate of the effect of each affirmative action regime. This experiment has not been conducted and will likely never be.

In lieu of the ideal experiment, we use remarkable data from several elite colleges and universities. The virtue of our approach is that our empirical model allows one to analyze alternative affirmative action regimes, holding all else constant, to investigate the effect on ex post efficiency. Thus, one can alter the selectiveness of colleges, the covariates observed, or the objectives pursued and analyze their effects on the relative efficiency of different affirmative action regimes.

Yet, although the data are rich, any nonexperimental analysis has important caveats. Potential selection problems arise due to (1) lack of data on students who the university did not admit and their counterfactual performance had they been accepted and (2) the information that admissions committees use is much richer than the covariates available to researchers. Thus, if we observe a student admitted with low SAT scores, it is likely that he wrote a stellar essay, had marvelous recommendation letters, and so on. For our purposes, the former is the most serious threat to the plausibility of our empirical estimates. Further, our simulations also assume that there are no spillover effects from attending colleges with higher mean quality. If the quality of a student is a function of their innate ability and the mean quality of her peers, our estimates will be biased.

The parameter that we are most interested in estimating is the inefficiency that a college can expect from practicing color-sighted versus color-blind admissions policies, relative to laissez-faire. The thought experiment is to imagine a college looking at a pool of applicants and comparing the set of students that they would hypothetically admit under our three policy regimes and the expected quality of the resulting classes. Yet, because we do not have data on actual applicant pools (using instead the set of admitted students as the virtual applicant pool and imagining colleges having to admit half of the students that they, in fact, admitted), the parameter we estimate—the expected quality of a student admitted in our simulations conditional on being in our sample—could differ substantially from what we intend. The usual remedies for selection such as estimating selection equations or reweighting data are not applicable here because we have no information on the population of interest.

Whether or not our estimates are reasonable depend on the conditional distribution of academic traits above and below a college's selection threshold. If, for example, the conditional distribution of academic traits is more similar in the applicant pool than our restricted sample and minorities in this sample have lower academic credentials, our estimates will be biased upward.[20] This is

20. If black academic credentials are superior to whites, our estimates biased downward.

likely the case if the white distribution has a larger right tail among accepted students. Conversely, if the conditional distribution is less similar in the applicant pool than our restricted sample and the minorities among the set of applicant who were not admitted are less academically distinguished than their white peers, our estimates will be biased downward.

We estimated the distributions of predicted college rank for each of our schools, by race, to help inform which way the selection may go. In all cases, the white conditional distribution (conditional on being in our sample) dominates the black distribution. Thus, if the conditional distribution (conditional on not being admitted) of blacks in the applicant pool—which we do not observe—is similar to whites, our forthcoming estimates of the short-run inefficiencies involved in practicing affirmative action are too large.

In summary, we recognize that selection is a potentially serious problem, but can offer no compelling remedy.

2.5 Results

Tables 1 and 2 present summary statistics for our sample of students in four liberal arts colleges and three research universities, respectively, broken down by institution and racial group. Black and Hispanic students score over 1 standard deviation below white and Asian students on the math and verbal sections of the SAT and have (on average) lower percentile ranks in high school, which is consistent with previous research. Among the socioeconomic variables, black students live in lower income zip codes with a substantial fraction of other blacks and have parents who are less likely to be college educated and dramatically less likely to be an alumnus of their child's college. A similar pattern holds for Hispanics—yet Asians outperform whites on the math section of the SAT, live in higher median income zip codes, and are more segregated from Blacks and Hispanics.

Table 3 reports results from the college-specific regression equations, which used academic and socioeconomic background variables to predict a student's class rank after 4 years of matriculation (equation (1)). The interpretation of the coefficients is standard: a one-unit change in the independent variable, all else constant, produces the reported change in the dependent variable. For instance, the coefficient on SAT math in College A is 4.04. This means that a 100-point increase in SAT scores is associated with a 4-point higher GPA rank at college graduation.

Interestingly, both a student's SAT verbal score and their high school rank are stronger predictors of college rank upon graduation than is his or her SAT math score. Parental education is also a strong predictor of college rank. After controlling for our three academic variables and parental education, the average income of a student's zip code is not statistically significant. Zip code racial demographics, however, are important predictors of college performance, even after controlling for other precollege characteristics. In six out of seven colleges, a student's college performance is positively (and significant) related to the fraction of Asians in that student's zip code and negatively

Table 1. Descriptive Statistics of Students in Four Liberal Arts Colleges, by Race, 1989

	College A				College B				College C				College D			
	Asian	Black	Hispanic	White	Asian	Black	Hispanic	White	Asian	Black	Hispanic	White	Asian	Black	Hispanic	White
College GPA rank	48.07	23.66	24.60	53.16	39.61	27.52	34.86	53.70	47.88	21.15	38.41	55.10	49.65	16.52	27.29	53.71
	(30.05)	(22.03)	(16.64)	(27.57)	(28.26)	(25.64)	(24.74)	(27.74)	(29.04)	(19.75)	(29.41)	(26.88)	(27.68)	(16.72)	(17.01)	(27.73)
SAT math	640.44	486.54	521.11	638.53	623.52	514.58	547.14	596.41	654.30	500.61	503.33	633.18	716.85	562.73	595.22	692.03
	(86.61)	(93.76)	(115.70)	(78.06)	(70.58)	(78.13)	(88.55)	(69.97)	(60.19)	(73.55)	(81.81)	(65.74)	(50.91)	(77.59)	(88.67)	(60.55)
SAT verbal	569.56	485.96	500.00	616.32	546.14	522.5	507.62	592.59	585.70	491.02	510.37	607.19	671.67	552.42	593.48	664.29
	(100.91)	(104.75)	(120.49)	(81.81)	(90.59)	(95.20)	(98.53)	(71.35)	(84)	(72.03)	(80.5)	(70.46)	(68.48)	(74.12)	(81.89)	(72.09)
HS percentile	88.48	76.49	76.43	89.62	87.93	84.49	92.12	89.96	92.72	82.04	87.54	93.20	96.84	84.09	94.13	94.63
	(8.50)	(18.18)	(14.60)	(9.62)	(9.64)	(16.86)	(6.35)	(9.52)	(7.65)	(15.73)	(9.63)	(6.57)	(3.35)	(14.72)	(5.29)	(7.10)
Mother college educated	0.64	0.46	0.56	0.80	0.55	0.46	0.05	0.58	0.71	0.51	0.44	0.81	0.56	0.48	0.17	0.78
	(0.48)	(0.51)	(0.51)	(0.40)	(0.50)	(0.51)	(0.22)	(0.49)	(0.46)	(0.51)	(0.51)	(0.39)	(0.50)	(0.51)	(0.38)	(0.42)
Father college educated	0.87	0.41	0.61	0.86	0.68	0.46	0.19	0.66	0.78	0.43	0.56	0.81	0.69	0.55	0.25	0.86
	(0.33)	(0.50)	(0.50)	(0.34)	(0.47)	(0.51)	(0.40)	(0.47)	(0.42)	(0.50)	(0.51)	(0.39)	(0.47)	(0.51)	(0.44)	(0.35)
Zip income	44,240	31,792	33,906	42,507	43,826	34,686	37,858	43,082	50,031	32,308	35,986	46,573	47,131	37,812	37,652	47,712
	(15,686)	(14,492)	(10,871)	(17,081)	(16,953)	(15,744)	(14,689)	(17,050)	(18,541)	(14,401)	(18,637)	(18,233)	(18,477)	(18,257)	(16,561)	(20,198)
Percent Asian in zip	0.10	0.02	0.04	0.03	0.10	0.02	0.11	0.03	0.08	0.03	0.04	0.03	0.08	0.04	0.05	0.03
	(0.16)	(0.03)	(0.03)	(0.04)	(0.14)	(0.01)	(0.09)	(0.05)	(0.10)	(0.04)	(0.03)	(0.05)	(0.09)	(0.05)	(0.05)	(0.04)
Percent black in zip	0.05	0.39	0.14	0.08	0.04	0.18	0.03	0.04	0.04	0.41	0.05	0.04	0.06	0.34	0.12	0.06
	(0.09)	(0.38)	(0.13)	(0.13)	(0.08)	(0.26)	(0.06)	(0.09)	(0.07)	(0.34)	(0.07)	(0.09)	(0.14)	(0.35)	(0.14)	(0.11)
Percent Hispanic in zip	0.05	0.05	0.19	0.05	0.09	0.09	0.34	0.04	0.06	0.09	0.22	0.04	0.08	0.08	0.24	0.04
	(0.06)	(0.09)	(0.10)	(0.09)	(0.13)	(0.14)	(0.32)	(0.07)	(0.10)	(0.14)	(0.31)	(0.07)	(0.12)	(0.12)	(0.21)	(0.06)
Legacy	0.01	0.02	0.06	0.05	0.02	0.04	0	0.10	0.07	0.06	0.11	0.14	0.06	0	0	0.15
	(0.10)	(0.13)	(0.24)	(0.22)	(0.15)	(0.20)	(0)	(0.30)	(0.26)	(0.24)	(0.32)	(0.35)	(0.23)	(0)	(0)	(0.35)
Male	0.40	0.47	0.67	0.46	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	0.54	0.45	0.29	0.55
	(0.49)	(0.50)	(0.49)	(0.50)	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	(0.50)	(0.51)	(0.46)	(0.50)
N	91	59	18	646	88	24	21	471	130	49	27	374	54	33	23	408

All data are drawn from the College and Beyond Database, except for the demographic data on zip codes, which were attained from the 1990 Census. The numbers in parentheses are standard deviations. See the data appendix for further details of the construction of the variables. HS, high school; n/a, not available.

Table 2. Descriptive Statistics of Students in Three Research Universities, by Race, 1989

	College E				College F				College G			
	Asian	Black	Hispanic	White	Asian	Black	Hispanic	White	Asian	Black	Hispanic	White
College GPA	59.45	25.44	28.84	52.17	51.92	20.16	41.70	53.06	55.92	33.02	47.87	51.03
rank	(25.07)	(20.68)	(21.24)	(28.36)	(27.47)	(20.51)	(27.37)	(27.69)	(28.71)	(27.09)	(29.36)	(28.25)
SAT math	731.18	612.44	641.6	706.54	690.47	541.78	623.57	651.41	614.43	515.05	569.25	605.48
	(48.33)	(65.31)	(73.61)	(59.41)	(60.69)	(81.80)	(74.25)	(73.17)	(101.11)	(73.55)	(91.24)	(80.83)
SAT verbal	648.49	595.24	598.60	650.69	578.22	523.24	550.00	583.63	549.65	502.20	508.88	557.17
	(86.16)	(65.78)	(69.55)	(69.51)	(89.37)	(82.59)	(86.37)	(70.13)	(90.59)	(99.62)	(83.91)	(80.85)
HS percentile	97.83	94.85	92.85	95.69	94.80	84.41	93.36	93.25	87.06	83.90	81.38	81.63
	(3.16)	(5.19)	(8.41)	(6.51)	(8.05)	(14.75)	(10.65)	(8.58)	(17.45)	(15.63)	(12.31)	(16.03)
Mother college	0.78	0.55	0.44	0.73	0.68	0.47	0.32	0.65	0.29	0.20	0.15	0.20
educated	(0.41)	(0.50)	(0.50)	(0.44)	(0.47)	(0.50)	(0.48)	(0.48)	(0.46)	(0.41)	(0.36)	(0.40)
Father college	0.93	0.49	0.48	0.84	0.87	0.45	0.58	0.79	0.32	0.17	0.35	0.24
educated	(0.25)	(0.50)	(0.50)	(0.37)	(0.33)	(0.50)	(0.50)	(0.41)	(0.47)	(0.38)	(0.48)	(0.43)
Zip income	51,773	39,858	41,374	50,741	45,006	33,223	33,077	45,470	35,565	26,042	35,578	42,848
	(18,964)	(20,783)	(17,922)	(20,460)	(16,146)	(14,568)	(11,168)	(18,715)	(16,959)	(11,383)	(20,796)	(21,054)
Percent Asian	0.07	0.03	0.07	0.04	0.08	0.03	0.04	0.03	0.06	0.02	0.03	0.03
in zip	(0.12)	(0.03)	(0.11)	(0.04)	(0.13)	(0.04)	(0.05)	(0.05)	(0.08)	(0.04)	(0.03)	(0.05)
Percent black	0.05	0.29	0.03	0.05	0.05	0.39	0.11	0.05	0.16	0.39	0.11	0.11
in zip	(0.10)	(0.34)	(0.07)	(0.11)	(0.09)	(0.33)	(0.19)	(0.10)	(0.18)	(0.28)	(0.18)	(0.17)
Percent	0.05	0.06	0.25	0.04	0.05	0.06	0.18	0.03	0.08	0.04	0.12	0.05
Hispanic	(0.07)	(0.07)	(0.29)	(0.07)	(0.06)	(0.11)	(0.23)	(0.06)	(0.13)	(0.08)	(0.18)	(0.08)
in zip												
Legacy	0.17	0.05	0.12	0.22	0.15	0.06	0.13	0.25	0.59	0.72	0.65	0.72
	(0.38)	(0.22)	(0.33)	(0.41)	(0.35)	(0.21)	(0.34)	(0.43)	(0.50)	(0.45)	(0.48)	(0.45)
Male	0.50	0.48	0.48	0.58	0.50	0.41	0.55	0.52	0.46	0.52	0.54	0.52
	(0.50)	(0.50)	(0.50)	(0.49)	(0.50)	(0.49)	(0.51)	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)
N	119	82	50	888	281	157	28	1353	68	98	48	1270

All data are drawn from the College and Beyond Database, except for the demographic data on zip codes, which were attained from the 1990 Census. The numbers in parentheses are standard deviations. See the data appendix for further details of the construction of the variables. HS, high school.

Table 3. Performance Equation: Predicted College Rank

	College A	College B	College C	College D	College E	College F	College G
SAT math	4.04 (1.39)	−0.60 (1.78)	5.08 (1.69)	7.57 (1.51)	6.33 (0.89)	1.61 (0.95)	9.41 (1.86)
SAT verbal	5.47 (1.31)	8.98 (1.62)	7.15 (1.68)	12.85 (1.30)	6.57 (0.89)	5.77 (0.95)	10.51 (1.64)
HS percentile	3.12 (1.11)	8.79 (1.42)	8.42 (1.84)	7.49 (1.79)	6.06 (0.54)	6.68 (0.79)	8.37 (1.69)
Mother college educated	2.58 (2.63)	8.40 (3.00)	−3.45 (3.76)	3.61 (2.15)	7.47 (3.40)	3.00 (1.46)	3.54 (2.65)
Father college educated	4.35 (2.99)	−3.76 (3.60)	6.07 (4.04)	5.48 (2.84)	5.52 (4.06)	4.34 (1.72)	3.46 (3.12)
Zip income	−0.04 (0.64)	−1.44 (0.80)	−0.47 (0.72)	−0.74 (0.44)	0.33 (0.40)	0.87 (0.38)	0.03 (0.61)
Legacy	4.66 (4.55)	0.59 (4.05)	0.65 (3.64)	−0.47 (1.96)	−0.39 (2.77)	−3.23 (1.54)	3.46 (3.43)
Percent Asian in zip	14.07 (16.83)	16.78 (16.82)	6.28 (19.45)	33.05 (13.58)	21.89 (13.47)	9.74 (9.39)	−8.53 (23.94)
Percent Black in zip	−11.72 (5.78)	−29.10 (10.99)	−14.26 (7.59)	−15.91 (5.31)	−9.29 (4.03)	−21.86 (4.03)	−13.92 (7.33)
Percent Hispanic in zip	−15.76 (11.21)	−22.15 (11.51)	−0.42 (11.40)	−3.24 (9.10)	0.83 (7.84)	−13.78 (9.01)	−26.31 (13.65)
Male	−4.77 (2.06)	n/a	n/a	−7.92 (1.66)	−10.07 (1.53)	−7.17 (1.34)	−2.16 (2.20)
R^2	0.16	0.21	0.19	0.37	0.28	0.16	0.26
Number of observation	761	429	512	494	955	1787	1419

College rank is percentiles in distribution of cumulative GPA among students who matriculated at that college in 1989. HS percentile is students' percentile in his high school. Mother and father's education are dummies for students' mother and father being college educated. Zip income is the median income of the student's zip code from the 1990 Census; n/a, not available. Increments: SAT variables 100 points, HS percentile 10 percentiles, ZIP income $10,000. We used dummies for the missing data. (Coefficients for these variables are not reported in this table.)

related to the fraction of Hispanics. All else equal, the fraction of blacks in a zip code is negatively related to college performance for all seven colleges. It is likely that these racial demographic correlations are capturing unobserved school or neighborhood quality factors that promote performance in college, but that are not captured in standardized tests. Perhaps the most surprising finding of Table 3 is that more than 60% of student variation in college rank remains unexplained at all of colleges, after taking account of students' pre-college characteristics.

Table 4 reports results from the auxiliary regressions that we imagine the colleges to have run if, when operating under a color blindness constraint, they needed to use academic and social background variables to forecast the likelihood that a student is black. As intuition might suggest, Blacks are more likely to live in own-race neighborhoods, reside in lower income zip codes, be non-legacies, and have lower scholastic achievement.

Tables 5–7 report the results of greatest interest, regarding the relative inefficiency of race-neutral alternatives (Table 5), the implication of such policies for the representation of various racial groups (Table 6), and the way that optimal color-blind affirmative action alters the weight given to various factors in the optimal admissions formula—that is, test scores, grades, and socioeconomic background measures (Table 7).[21] Bear in mind that we measure the performance of a policy in terms of the average of the class rank predicted for the students admitted under that policy.

Table 5 reports our estimates of the relative performance of conventional (i.e., color sighted) affirmative action and four alternatives policies (laissez-faire performance at each college has been normalized at 100). The columns in the table represent institutions of higher education. The first four (A–D) are liberal arts colleges, and the last three (E–G) are research universities. The rows represent five different policy regimes: random admissions; LF policies when colleges are artificially prohibited from knowing students' SAT scores and high school grades, respectively; color-sighted affirmative action; and color-blind affirmative action. As theory predicts, in every case color-blind policies perform less well than do color-sighted policies. The magnitude of efficiency loss from employing color-blind rather than color-sighted affirmative action varies across institutions, ranging from less than 1 percentage point at College B to more than 6.6 percentage points at College D (a small, elite institution in the northern portion of the United States). Notice that at most colleges, the loss of selection efficiency associated with going from color-sightedness to color blindness (given the same representation target) is comparable to, and sometimes even exceeds, the loss of efficiency that would arise

21. Both LF and CS optimal admissions policies use the same weights (those derived from the regression predicting college class rank), whereas the CB policy employs weights that are "biased" in order to exploit the fact that some variables are more closely correlated (positively or negatively) than are others with a student's being black.

Table 4. Race Equation: Probability of Being Black

	College A	College B	College C	College D	College E	College F	College G
SAT math	−0.06 (0.01)	−0.06 (0.01)	−0.07 (0.01)	−0.10 (0.02)	−0.10 (0.01)	−0.08 (0.01)	−0.06 (0.01)
SAT verbal	−0.02 (0.01)	−0.03 (0.01)	−0.06 (0.01)	−0.04 (0.01)	−0.03 (0.01)	−0.02 (0.01)	−0.03 (0.01)
HS percentile	−0.04 (0.01)	0.01 (0.01)	−0.07 (0.01)	−0.05 (0.01)	0.01 (0.01)	−0.04 (0.01)	0.01 (0)
Mother college educated	0.02 (0.02)	0.01 (0.02)	0.03 (0.03)	0.01 (0.02)	0.04 (0.02)	−0 (0.01)	−0.02 (0.03)
Father college educated	−0.11 (0.02)	−0.04 (0.02)	−0.08 (0.03)	−0.02 (0.03)	−0.09 (0.02)	−0.05 (0.01)	0.02 (0.04)
Zip income	0.01 (0)	0.01 (0.01)	0 (0.01)	0.01 (0.01)	0.01 (0)	0 (0)	0 (0)
Legacy	0.01 (0.03)	−0.02 (0.03)	−0.01 (0.03)	−0.09 (0.03)	−0.06 (0.02)	−0.05 (0.01)	−0.04 (0.02)
Percent Asian in zip	−0.25 (0.12)	−0.06 (0.11)	−0.10 (0.15)	0.14 (0.20)	0.00 (0.11)	−0.01 (0.08)	0.11 (0.12)
Percent Black in zip	0.57 (0.04)	0.57 (0.07)	0.85 (0.06)	0.57 (0.06)	0.64 (0.04)	0.78 (0.03)	0.44 (0.04)
Percent Hispanic in zip	−0.02 (0.08)	−0.02 (0.08)	−0.08 (0.09)	−0.13 (0.12)	−0.15 (0.08)	−0.00 (0.07)	0.05 (0.07)
Male	0.01 (0.02)	n/a	n/a	0.01 (0.02)	0.02 (0.01)	0.01 (0.01)	0.04 (0.01)
R^2	0.41	0.22	0.55	0.37	0.39	0.42	0.21
Number of observations	761	429	512	494	955	1787	1419

Dependent variable is student's probability of being black. HS percentile is students' percentile in his high school. Mother's and father's education are dummies for students mother and father being college educated. Zip income is the average income of the student's zip code from the 1990 Census; n/a, not available. Increments: SAT variables 100 points, HS percentile 10 percentiles, ZIP income $10,000. We used dummies for the missing data. (Coefficients for these variables are not reported in this table.)

Table 5. Relative Performances of Color-Blind and Color-Sighted Policies, by Race Constraint

	College A	College B	College C	College D	College E	College F	College G	Average
Random admissions	86.29	87.22	80.43	85.04	80.67	85.62	81.26	83.62
Laissez-faire without SAT	96.96	98.50	97.00	93.90	95.01	98.45	96.65	96.88
Laissez-faire without HS percentile	99.77	98.02	97.89	99.58	99.56	98.67	97.24	98.57
Color sighted	97.97	99.55	97.64	97.43	98.66	98.59	99.77	98.68
Color blind	94.28	98.67	95.33	90.82	96.40	95.95	98.74	96.16

Predicted college rank of a student is estimated by the OLS regression. For each policy, we compute the average predicted college rank of the admitted class. We call this value the performance of the policy. To compute the relative performance, we index laissez-faire's performance as 100. For example, color-sighted relative performance = (color-sighted performance × 100)/(laissez-faire performance). Average is the population-weighted average.

if colleges had no interest in racial representation, but were constrained from using students' grades or test scores in the admissions process. (That random admissions costs no college more than 20% in efficiency is clearly due to the fact that our sample consists of matriculates, not applicants: grossly

Table 6. Diversity of Entering Class Under Alternative Policies

	College A	College B	College C	College D	College E	College F	College G
Laissez-faire							
Asians	34	30	62	26	74	174	41
Blacks	2	5	1	0	4	11	26
Hispanics	2	3	5	1	7	10	14
Whites	376	265	223	232	484	734	662
Color sighted							
Asians	29	30	53	24	73	165	38
Blacks	29	12	24	16	41	78	49
Hispanics	2	3	5	1	5	9	16
Whites	354	257	208	217	449	676	639
Color blind							
Asians	32	28	52	25	71	165	43
Blacks	27	12	26	17	43	73	51
Hispanics	5	4	7	4	8	10	16
Whites	349	258	205	213	447	680	633

This table shows the effects of laissez-faire, color-sighted, and color-blind policies on the composition of the admitted class. Note that color-blind policy is targeted for black students only. The number of students are not integers under color-blind policy. This is because the college assigns an admission probability for two students strictly between zero and one. The college has a linear programming problem: to maximize a linear function with respect to linear constraints. The college assigns a value to each student (with respect to his expected college rank and probability of being black) and admits the students with higher values. However, on the margin, the college is indifferent between two students (two students have equal values). That is why admitting a fraction of these students is consistent with the result being optimal.

Table 7. Weight on Students' Characteristics in the Admission Formula for Laissez-Faire and Color-Blind Policies, by Representation Goal

	SAT math		SAT verbal		HS percent		Mother educated		Father educated		Income		Percent black		Percent Asian		Percent Hispanic	
	LF	CB	LF	CB	LF	CB	LF	CB	LF	CB	LF	CB	LF	CB	LF	CB	LF	CB
College A	4.04	0.16	5.47	4.18	3.12	0.53	2.58	3.87	4.35	−2.76	−0.04	0.61	−11.72	25.14	14.07	−2.10	−15.76	−17.05
College B	−0.60	−4.88	8.98	6.84	8.79	9.50	8.40	9.11	−3.76	−6.61	−1.44	−0.73	−29.10	11.56	16.78	12.50	−22.15	−23.58
College C	5.08	1.68	7.15	4.24	8.42	5.02	−3.45	−1.99	6.07	2.18	−0.47	−0.47	−14.26	27.02	6.28	1.42	−0.42	−4.31
College D	7.57	−3.74	12.85	8.33	7.49	1.83	3.61	4.74	5.48	3.22	−0.74	0.39	−15.91	11.77	33.05	48.89	−3.24	−17.94
College E	6.33	−0.47	6.57	4.53	6.06	6.74	7.47	10.19	5.52	−0.60	0.33	1.01	−9.29	34.2	21.89	21.89	0.83	−9.36
College F	1.61	−1.96	5.77	4.88	6.68	4.89	3.00	3.00	4.34	2.11	0.87	0.87	−21.86	12.98	9.74	9.29	−13.78	−13.78
College G	9.41	6.38	10.51	9.00	8.37	8.87	3.54	2.53	3.46	4.47	0.03	0.03	−13.92	8.30	−8.53	−2.98	−26.31	−23.79

"underqualified" individuals are unlikely to have been admitted to these selective institutions and, therefore, are significantly underrepresented in our sample.)

Table 6 shows the consequence of our three policy regimes for the overall ethnic/racial composition of the admitted class. Because color-blind affirmative action shifts weight from academic characteristics to social characteristics, such policies directed toward blacks will concurrently help Hispanics and low-income whites, whereas color-sighted affirmative action will not. Table 6 also sheds light on the inefficiencies described above. Consider College D. Under an LF regime, the college admits no black students. This implies that all blacks are in the left tail of the academic trait distribution, conditional on being accepted. But, affirmative action requires College D to admit 16 black students, which explains the efficiency hit. Conversely, College B would admit five black students on their own accord and are required to admit only seven more due to affirmative action. As such, College B has very little loss in efficiency in the short-run.

Table 7 reports our calculations of the weights on students' characteristics in the admission formula that are employed under optimal LF and CB policies. In effect, colleges are assigning a score to each student and admitting that half of the applicant pool with higher scores. The numbers in Table 5 are simply the coefficients used in a linear formula to derive a student's score from that student's academic and socioeconomic traits. It is clear from the illustrative empirical results reported in Table 7 that optimal CB admissions policy gives less weight to test scores, more weight to high school grades, and more weight to social background factors than does optimal policy under the LF-CS regimes. As such, it is not surprising that CB admissions policy targeted on blacks tends also to raise admissions rates for Hispanics while lowering them for whites and Asians, as Table 6 reveals. This is the crucial point in which these estimates and the theoretical model developed in the previous section interact: The short-run efficiency of color-blind affirmative action may not seem so great (and it would shrink toward zero if we could add more variables that are correlated with race.) Yet, because color-blind policy shifts weight from performance-related traits to social characteristics that are weakly correlated with achievement, it lowers incentives for applicants to invest in the traits valued by selectors.

Throughout the above analysis, we made two assumptions that warrant further emphasis. First, we imagined that the colleges could admit only half of their original matriculating classes. To gauge the importance of that assumption, Figure 2 plots the relative efficiency of various affirmative action policies at each college while allowing colleges to admit a proportion of their original classes that ranges from 10% to 90%. One can see from the figure that the tighter is the capacity constraint (i.e., the lower the "percent admitted"), the more efficient is CS policy relative to random admissions. Yet, the inefficiency of CB relative to CS policy is essentially independent of the percent admitted at all institutions except College D. We conclude that our results are not sensitive to this first assumption.

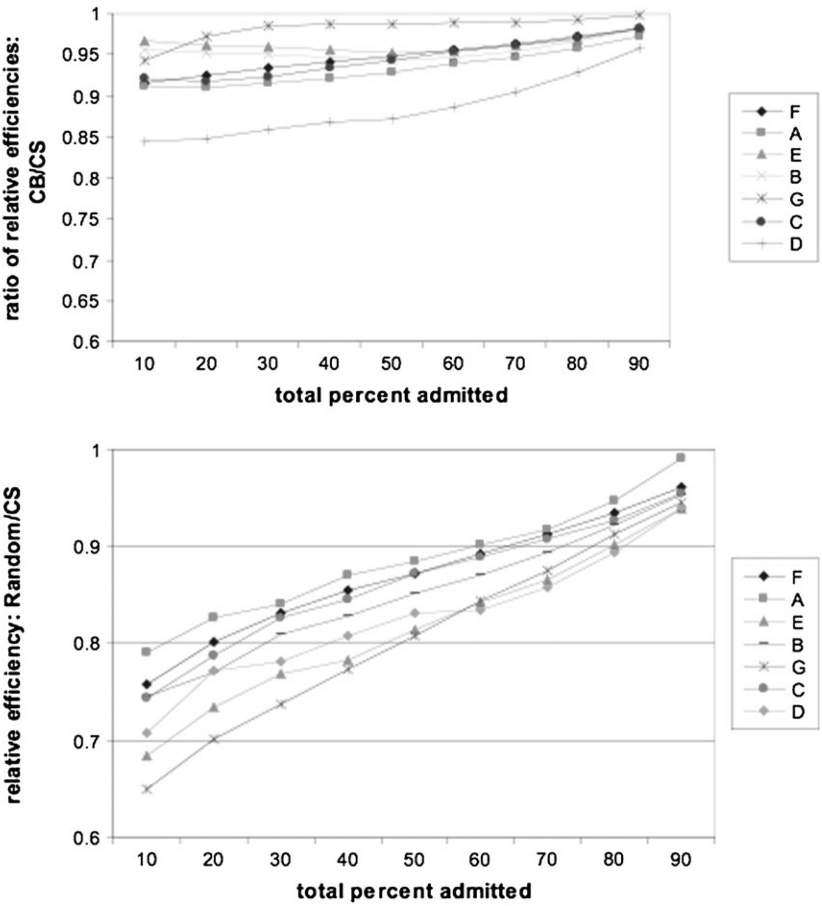

Figure 2. Ratios of Relative Efficiencies.

Second, we assumed that colleges strive to achieve the same level of diversity under our hypothetical color-blind constraints as they had actually achieved under what we must presume to have been a color-sighted system. But, this assumption is implausible: if imposing blindness raises the cost of affirmative action, it stands to reason that colleges would then consume less of it! Further, as Sander (1997) demonstrates, the marginal cost to academic goals of racial affirmative action can rise rapidly with the percent minority students being admitted. Given this finding, it is plausible that the relative efficiency of color-blind affirmative action might also vary with the magnitude of the racial representation target. Figure 3 tries to detect such nonlinearities in our data. The x axis measures the affirmative action goal of each school. The y axis measures the efficiency of color-blind affirmative action relative to laissez-faire. Interestingly, as Sander (1997) observed in a different context,

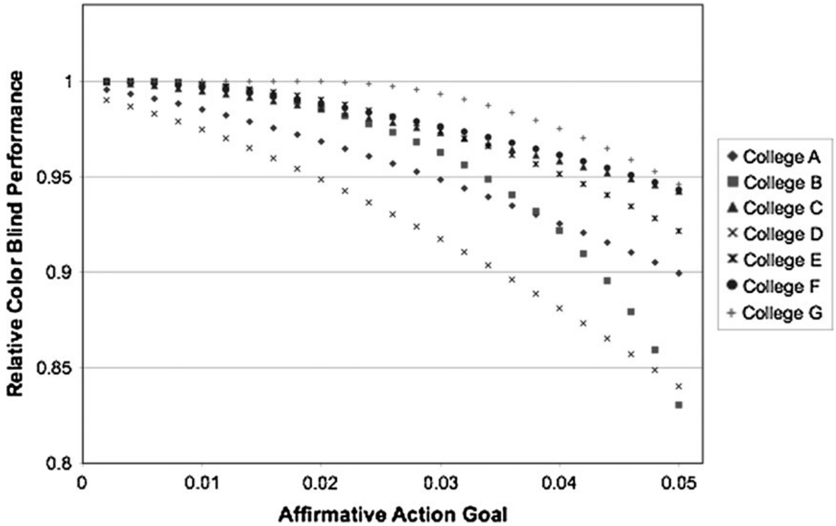

Figure 3. The Trade-off Between Efficiency and Representation.

we find a distinct nonlinearity involved in most schools between representation and efficiency. (Notice that College D, where color blindness was most inefficient, also has the most stark nonlinear relationship.)

3. The Long-Run Consequences of Color-Blind Affirmative Action

There are several reasons to expect that color-blind policies may undermine the efficiency of the selection process in the long-run. Unfortunately, with the current data it is not possible for us to empirically estimate the potential magnitude of these long-run effects.[22] However, by formulating a rigorous theoretical model of the selection problem with endogenous traits, we can gain some insight into the main issues.[23]

Imagine, then, that a continuum of applicants (students) of unit mass consists of two racial groups, $R \in \{1, 2\}$, where $\lambda \in (0, 1)$ is the proportion

22. One could try to estimate the impact of the percentage plans discussed in the introduction on student effort, using data across states over time. Given a measure of effort (e.g., average SAT scores or school attendance rates) in various states over several years, a difference-in-differences model could be estimated by contrasting student effort before and after the implementation of affirmative action, as between states using blind versus those using sighted policies. Yet, since the adoption of blind rather than sighted policies is clearly not exogenous at the state level, it is not clear what would be learned by such an exercise.

23. Our model implicitly assumes (1) that the law permits the pursuit of racial diversity with policies that are not explicitly contingent on an applicant's race and (2) that perfect proxies for race do not exist (i.e., there are no variables which, taken together, are perfectly correlated with race while being unrelated to college performance).

belonging to group 1. Applicants seek to be accepted by any one of a large but finite number, N, of identical firms (colleges). Each applicant is randomly assigned to a firm, and so each firm faces an applicant pool of measure $1/N$ that is the statistical replica of the overall population. Let $\omega > 0$ be the gross value to an applicant of being accepted. Each firm can accept at most the fraction $c \in (0, 1)$ of those who apply. Firms prefer to accept the better qualified applicants and take the distribution of characteristics in their applicant pools as given, independent of their acceptance policies.

Prior to being assigned, applicants make an ex ante binary effort decision $e \in \{0, 1\}$ that affects their qualifications ex post. The incentive effects of affirmative action will be reflected in this model by the way that alternative policies alter the distribution in the student population of this binary effort variable and the resulting distribution of qualifications. We assume that low effort ($e = 0$) is costless, but high effort ($e = 1$) entails a cost, $k \geq 0$, for an applicant. The frequency distribution of effort cost differs between racial groups. Let $G_R(k)$ be the fraction of group R with effort cost less than or equal to k, and let $g_R(k)$ be the associated density function. Let $G(k) = \lambda G_1(k) + (1 - \lambda)G_2(k)$ be the cumulative distributive function (CDF) of effort cost for the overall population, with $g(k)$ being the associated population density function. We assume group 2 is disadvantaged in the sense that it has a uniformly less favorable cost distribution than group 1: $[g_1(c)/g_2(c)]$ is a monotonically decreasing function of c.

An applicant's qualification (as perceived and valued by firms) is a stochastic function of effort.[24] Let t be a number representing an applicant's qualification, let $F_e(t)$ be the probability that effort e leads to a level of qualification less than or equal to t, and let $f_e(t)$ be the associated density function. High effort is assumed to increase an applicant's qualification in the following sense: $[f_1(t)/f_0(t)]$ is a monotonically increasing function of t.

Finally, let π_R represent the fraction of applicants in group R who choose action $e = 1$ ($R = 1, 2$), with π being the fraction of all applicants who exert high effort. The variables π_R are endogenous and will depend on incentives for applicants to take high effort created by the firms' acceptance policies. In a population where the fraction π exerts high effort, the CDF of the distribution of qualifications is denoted $F(\pi, t)$; $f(\pi, t)$ denotes the associated density function.

An applicant is characterized by his or her racial group and degree of qualification. A firm's *acceptance policy* must be some function $A(R, t)$ representing the probability that an applicant of racial group R and qualification t is accepted. A firm's policy is color blind if $A(1, t) = A(2, t)$ for almost every t.

24. We emphasize that no asymmetry of information between firms and applicants is being assumed here. Qualifications are perfectly and costlessly observable by firms. Our assumption is that, when an applicant chooses effort ex ante, the extent of qualification that results ex post is random at the individual level. Because there is a continuum of applicants, the distribution of qualifications in any population depends only on the fraction of applicants in that population who have chosen high effort. We assume that firms care only about an applicant's qualifications ex post and not about the effort taken ex ante.

We consider the behavior of firms under three possible policy regimes: laissez-faire, color-sighted and color-blind affirmative action (LF, CS, and CB, respectively). In each case, we assume that firms take the proportion of high-effort applicants in each racial group as given when deciding upon an acceptance policy.

Under LF, firms are unconcerned with diversity so they ignore group identity information. Thus, given the proportion of high-effort applicants, firms choose an acceptance policy which maximizes the expected quality of those admitted, subject to their capacity constraint. The best LF policy is a color-blind threshold policy, where firms accept the fraction c of the applicant pool with the highest qualifications.

Let r_2^* be the acceptance rate for group 2 that is obtained under this LF optimal policy. We formalize affirmative action (either the color-blind or the color-sighted variety) by positing that firms seek an acceptance rate for group 2 members, r_2, that exceeds r_2^*, but is no greater than that implied by population parity. Given their beliefs about the fraction in each group of applicants who have chosen high effort, we require firms to choose an acceptance policy under which they anticipate to accept group 2 applicants at the rate, r_2. The aggressiveness of the affirmative action policy pursued by firms is taken to be exogenous throughout this analysis. In light of the capacity constraint and given this two-group setup, a representation target for group 2, r_2, necessarily implies a target for group 1, r_1.

Consider now a firm's selection problem under a CS policy regime with representation target r_2. Taking the fraction of high-effort applicants in each group as given, firms choose an acceptance policy for each group which maximizes the expected quality of those admitted, subject to their capacity constraint and ensuring that group 2 applicants are adequately represented among those who are selected. Under color-sighted affirmative action, firms follow distinct threshold policies for each group, accepting the fraction r_R of the group R applicant pool with the highest qualifications, $R = 1, 2$.

Lastly, consider the firms' behavior in the CB regime, with the representation target for group 2 given. Firms again take the distribution of high effort in both groups as given, but now they must choose a color-blind acceptance policy function so as to maximize mean qualifications of those accepted while anticipating to generate the desired representation of group 2 members. The key issue is that a firm's problem cannot be a function of racial identity. This problem is a linear program in an infinite-dimensional space. Such programs are studied extensively in Anderson and Nash (1987), though our technical appendix shows that we need not solve this problem explicitly to characterize the *equilibrium* distribution of applicant qualifications under a CB policy regime in this model.

Imposing the color-blind constraint on firms that remain intent on achieving more representation for the disadvantaged group than occurs under laissez-faire must lead to a situation in which some applicants are accepted, whereas others with higher qualifications are rejected, thereby undercutting applicants' incentives to exert high effort. This is a general feature of color-blind affirmative action policies, and it is the basic reason that such policies must be inefficient over the longer run, relative to the color-sighted alternative.[25]

Regardless of the regulatory regime, an applicant with exogenous characteristics (R, k) who anticipates firms to employ the acceptance policy $A(R, t)$ will exert high effort only if the costs of doing so are no greater than the benefit.

An equilibrium in this model is an acceptance policy for firms and an effort supply function for applicants, which are mutual best responses to one another (see Proposition 1, Appendix).

3.1 The Impact of Affirmative Action on Applicant Qualifications in Equilibrium

To describe how affirmative action policies affect the equilibrium distribution of qualifications among applicants in the two groups in our model we must introduce some additional terminology. If in LF equilibrium the marginal applicant has low (high) qualifications, then we will say that acceptance standards are "loose" ("tight"). Furthermore, if the representation target $r_2 \approx r_2^*(c)$ $(r_2 \approx c)$ then we will say that the affirmative action goal is "weak" ("strong").

Proposition 2 in the Appendix establishes the following set of results. (1) If standards are loose in LF equilibrium, then the pursuit of sufficiently weak affirmative action goals with CS policies increases qualifications among the advantaged group and decreases qualifications among the disadvantaged group, thereby widening the racial qualifications gap. (2) If standards are tight in LF equilibrium, then weak CS affirmative action decreases qualifications among the advantaged and increases qualifications among the disadvantaged, thereby narrowing the racial qualifications gap. (3) If laissez-faire equilibrium standards are neither loose nor tight, then sufficiently strong CS affirmative action goals must decrease the qualifications of both groups.

Finally, Proposition 3 considers the effect of color-blind affirmative action on applicant qualifications in equilibrium. To state the result we need one last

25. This property would not hold if, in the manner of Chan and Eyster (2003), we were to impose some kind of monotonicity constraint on firms (e.g., requiring that $A(t)$ be nondecreasing, out of the incentive compatibility concern that applicants not see any gain from underreporting their qualifications.) Still, the basic point we are making here would remain valid, even if we were to impose monotonicity. Under such a constraint, the firm's problem can be reformulated so that it becomes (the dual of) what Anderson and Nash (1987, Section 4.4) call a "continuous semi-infinite linear program." If we apply their Theorem 4.8 (page 76) to this reformulated problem, we can conclude that with a monotonicity constraint the firm's optimal acceptance policy can be expressed as a step function with at most two points of discontinuity. This, in turn, implies that there will be levels of qualification t and s, with $t < s$, such that $A^*(t) > 0$ and $A^*(s) < 1$. That is, some applicants are accepted with a probability strictly greater than zero, whereas others with higher qualifications are accepted with a probability strictly less than one, again undercutting applicants' incentives to exert high effort.

But, the main point we wish to emphasize is that, once applicants' qualifications are allowed to be endogenous in the manner that we follow here, the imposition of such a monotonicity constraint on the firm's acceptance policy is irrelevant for determining the distribution of qualifications in equilibrium. This is because (as we show in the Appendix), given the capacity and representation constraints, all feasible color-blind affirmative action policies for firms generate the same (diminished) effort incentives for applicants.

definition: We will say that a selection problem is characterized by "elitism" if no feasible acceptance policy by firms can induce "high-cost" applicants to exert high effort. Now, suppose the condition of elitism obtains. Then we show in the Appendix that (1) for every affirmative action goal r_2, there is a unique color-blind affirmative action equilibrium; (2) as r_2 rises, the level of qualifications in both groups declines, as does the qualifications gap between the groups; (3) in the limit, as r_2 approaches c (population parity as a goal), the proportion of applicants choosing high effort in the color-blind affirmative action equilibrium approaches zero; however, (4) there does not exist an equilibrium in pure strategies under color-blind affirmative action that implements the representation target of population proportionality, $r_2 = c$.

This last is a stark result that warrants further emphasis.[26] Consider an extreme example: one way to achieve population proportionality for all groups is to select from among candidates for a limited number of positions at random, with every applicant facing the same chance of success. This assures (with large numbers of applicants and statistical independence of applicant traits) that the fraction of successful candidates from any group equals the fraction of applicants from that group. Yet, random selection gives applicants no incentive to acquire traits valued by the selector. In equilibrium, the population of applicants (from all groups) will be much less distinguished under random selection, despite the fact that those selected will indeed be racially diverse. Appendix Proposition 3 demonstrates that the intuition of this extreme example extends to the general case.[27]

Our principle theoretical conclusion is that color-blind affirmative action entails a basic trade-off between incentives and representational goals. If firms are constrained to be color blind but continue to value diversity, they will act in such a way as to "flatten" the function that relates a worker's probability of being accepted (in equilibrium) to that worker's level of qualification: Some lower qualification workers must have a greater chance of being accepted under color-blind affirmative action, and some higher qualification workers must have a smaller chance. (Otherwise, the disadvantaged group, which has relatively more low-qualification members, cannot have its representation increased.) This flattening of the link between qualifications and success undercuts incentives for all workers to exert preparatory effort by reducing the net benefit of investment.

Beyond the narrow definitions of efficiency employed in this article, there are at least two scenarios where color-blind affirmative action might be

26. We have further results along these lines in a more general setting. See the article by Fryer and Loury (2003), which considers the problem of equilibrium and optimal handicapping (i.e., affirmative action or, more generally, "categorical redistribution") in winner-take-all markets.

27. There is an interesting externality here that promotes long-run inefficiencies. Individual selectors drawing on a large, common pool of prospective applicants may not take into account that their choice of selection criteria alters the distribution of traits in the overall applicant population. This makes the adoption of a random selection method look like a low-cost move for any given selector. But when all selectors make this choice, they are all worse-off than they would be if none of them made it.

preferred to color-sighted policies. The most obvious is that color-blind approaches may be more viable politically. Diversity-promoting policies not explicitly contingent on an applicant's race seem, in the current climate, either to fly below the legal radar screen or simply to be less objectionable than explicitly racial policies.[28] Second, to the extent that one is concerned with the relative reputation of minorities within firms, color blindness may be preferred since, although a "blind" selection mechanism reduces the reputation of the average selected individual relative to a "sighted" policy, it increases the relative reputation within firms of selected members from the preferred group.

4. Beyond Race

The issues explored in this article are of more general interest, beyond the study of racial equity.[29] There are many contexts in which a firm or public authority distributes some resource across a heterogeneous, categorically diverse population, with the dual objectives of allocating that resource to the most productive members of the population while avoiding an undue categorical disparity in receipt of the benefit.[30] A state government may need to distribute funds for public works among competing cities and towns, aiming to allocate the funds where they are most needed (or can best be made use of), while limiting any resulting disparity among jurisdictions. Similarly, a supplier of consumer credit (or insurance) may need to screen applicants according to creditworthiness (or insurability), without thereby generating a customer base with too few racial minorities. When the observable individual traits that are positively associated with creditworthiness (or insurability) are less frequently present in one population group than another, then simply screening out the least qualified applicants could lead to a stark disparity in rates of selection between groups. For political, economic, or legal reasons, such an outcome might be undesirable. However, it may also be undesirable in such settings to explicitly discriminate among applicants based on (race or sex defined) group identities. This situation leads to the posing of an analytical problem nearly identical to the one investigated in this article.

Alternatively, consider a customs union—for example, the European Common Market. Imagine that a member state wants to favor its domestic producers of some good, but cannot do so directly without violating the trade agreement. Imagine further that members of the customs union are permitted to impose quality standard regulations, which all goods, no matter where they originate, must meet. For instance, some Germans may want to limit imports into their country of Dutch beer, but may be forbidden to bar such products by

28. Justice O'Connor in Croson (Richmond v. J. A. Croson Co., 488 U.S. 469) and Adarandand (Adarand Constructors v. Pena, 515 U.S. 200) can be read as affirming the latter view.

29. There is a clear relation between our analysis of color-blind affirmative action and policy targeting in international trade. See Bhagwati (1971) for a nice survey. We are grateful to Avinash Dixit for pointing us to this literature.

30. Akerlof (1978) investigates a related problem of "tagging" in the context of optimal taxation and welfare programs.

Common Market rules. Still, they can require of any beer sold in Germany that it have so much hops, so little preservative, come in kegs that are made of a particular wood, and so on.[31]

More generally, let the country have some preferences about what its quality standards should be and suppose that the relative costs to domestic and foreign producers of meeting different standards are known. Suppose quality has two dimensions and that, compared to foreign producers, domestic firms are at an absolute cost disadvantage when forced to produce the laissez-faire optimal quality vector. So, domestic producers would get a relatively low market share under the laissez-faire optimal (i.e., disinterested) quality regulations. However, suppose domestic firms have a *comparative* cost advantage over foreigners in satisfying one dimension of quality. Then, by biasing regulation so as to give greater importance to that dimension of quality, the country in question can raise its domestic firms' market share without appearing to practice protectionism, but at the expense of having a less than optimal (given their natural preferences for quality) set of regulations. Again, we have arrived at a formulation analogous to the model studied here.

Appendix: Technical Proofs

In this section, we provide a formal derivation of the model and results in Section 3.

Firm Behavior

An applicant is characterized by the pair (R, t). A firm's acceptance policy must be some function $A(R, t)$ representing the probability that an applicant with characteristics (R, t) is accepted.

Definition 1. Firm's policy is color blind if $A(1, t) = A(2, t)$ for almost every t.

Under LF: Firms choose the function $\{A(t)\colon 0 \le t \le 1\}$, with $0 \le A(t) \le 1$, so as to:

$$\max\left\{\int tA(t)f(\pi, t)\mathrm{d}t\right\}, \quad \text{subject to} \int A(t)f(\pi, t)\mathrm{d}t = c.$$

The solution to this problem is given by: $\{A^*(t) = 1,\ t \ge t^*$, and $A^*(t) = 0,\ t < t^*\}$, where $F(\pi, t^*) = 1 - c$.

31. Indeed, there is a real case involving Beck (German) and Heineken (Dutch) beers, in which the European Court of Justice prohibited Germany from enforcing its purity requirements for beer against beverages imported from other members of the European Commission (see Commission of the European Communities v Federal Republic of Germany, Case 178/84, Judgment of 12 March 1987, 1987 ECR 1227). We are grateful to William James Adams for bringing this example to our attention.

Consider now a firm's selection problem under a color-sighted affirmative action policy regime, with representation target r_2. Taking (π_1, π_2) as given, firms choose $\{A(R, t)$: $0 \le t \le 1\}$, with $0 \le A(R, t) \le 1$, so as to:

$$\max\left\{\int tA(R,t)f(\pi_R,t)\mathrm{d}t\right\},\quad \text{subject to} \int A(R,t)f(\pi_R,t)\mathrm{d}t = r_R,\ R \in \{1,2\}.$$

Here the problem is solved separately for each group. As before, it is clear that the solution to this problem involves $\{A^*(R, t) = 1,\ t \ge t_R^*$, and $A^*(R, t) = 0$, $t < t_R^*\}$, where $F(\pi_R, t_R^*) = 1 - r_R$, $R = 1, 2$. Thus, under color-sighted affirmative action firms follow distinct threshold policies for each group, accepting the fraction r_R of the group R applicant pool with the highest qualifications, $R = 1, 2$.

Consider the firms' behavior in the CB regime, with the representation target for group 2 given as $r_2 < c$. Firms again take (π_1, π_2) as given, but now must choose a color-blind acceptance policy function so as to maximize mean qualifications of those accepted while anticipating to generate the desired representation of group 2 members. So, the problem for firms becomes choosing $A(t)$ so as to:

$$\max\left\{\int tA(t)f(\pi,t)\mathrm{d}t\right\},\quad \text{subject to} \int A(t)f(\pi_R,t)\mathrm{d}t = r_R,\quad R = 1,2.$$

An equivalent way of expressing the firm's problem under a CB regime is as follows:

$$\max\left\{\int tA(t)f(\pi,t)\mathrm{d}t\right\},\quad \text{subject to} \int [A(t)f(\pi,t)\mathrm{d}t] = c \ \text{ and}$$
$$\int [A(t)f(\pi,t)\xi(\pi_1,\pi_2,t)\mathrm{d}t] = (1-\lambda)r_2,$$

where

$$\xi(\pi_1,\pi_2,t) = \frac{(1-\lambda)f(\pi_2,t)}{f(\pi,t)}.$$

Thus, $\xi(\pi_1, \pi_2, t)$ is the conditional probability that an applicant belongs to the group 2, given group-specific high effort rates (π_1, π_2), and given that the applicant's level of qualification is t. This problem is a linear program in an infinite-dimensional space, as both the objective and the constraints may be regarded as linear functionals of the infinite-dimensional control variable $\{A(t)$: $t \in [0, 1]\}$. Such programs are studied extensively in Anderson and Nash (1987), though we will see momentarily that we need not solve this problem explicitly to characterize the equilibrium distribution of applicant qualifications under a CB policy regime in this model.

Let μ' be the multiplier on the capacity constraint and let θ be the multiplier on the representation constraint in the firm's CB optimization problem stated above. Then, by the infinite-dimensional analogue of the Kuhn-Tucker

Theorem, the solution to this problem is given by $\{A^*(t) = 1, t + \theta\xi(\pi_1, \pi_2, t) \geq \mu'$; and $A^*(t) = 0, t + \theta\xi(\pi_1, \pi_2, t) < \mu'\}$. Thus, an optimal acceptance policy function has the property that, for almost every t, either $A^*(t) = 1$, or $A^*(t) = 0$. It is also clear that, since $r_2^* < r_2$, this optimal policy under color-blind affirmative action cannot be a threshold policy. Hence, we may conclude that there will be levels of qualification t and s, with $t < s$, such that $A^*(t) = 1$ and $A^*(s) = 0$.

Observe that any color-blind acceptance policy function, $A(t)$, defines two numbers: $A_e = \left\{\int A(t) f_e(t)\mathrm{d}t\right\}, e \in \{0, 1\}$. A_e is the probability that an applicant who exerts effort e is accepted, and $\omega[A_1 - A_0]$ is the expected benefit to an applicant from choosing action $e = 1$ instead of $e = 0$. So, all acceptance policies that generate a common value for $[A_0 - A_1]$ must induce the same behavioral response from applicants. Now, given that a firm's representation target for group 2 is r_2 and given that its capacity implies $c = \lambda r_1 + (1 - \lambda) r_2$, it follows that the firm's constraints under a CB policy regime can be written as follows:

$$\pi_R A_1 + (1 - \pi_R) A_0 = r_R, \quad R = 1, 2.$$

From these equations it follows that $(A_1 - A_0) = \frac{[c - r_2]}{[\lambda(\pi_1 - \pi_2)]}$.

Applicant Behavior

Regardless of the regulatory regime, an applicant with exogenous characteristics (R, k) who anticipates firms to employ the acceptance policy $A(R, t)$ will behave in accordance with the effort supply function, $e^*(R, k)$, given as follows:

$$e^*(R, k) = 1 \quad \text{if } \omega \int [A(R, t)[f_1(t) - f_0(t)]\mathrm{d}t] \geq k, \quad \text{otherwise} \quad e^*(R, k) = 0.$$

That is, applicants choose to exert high effort if and only if their effort costs are below some threshold. Therefore, anticipating the acceptance policy function, $A(R, t)$, the proportion of group R who end up choosing the action $e = 1$ is given simply by:

$$\pi_R^* = G_R\left(\omega \int A(R, t)[f_1(t) - f_0(t)]\mathrm{d}t\right), \quad R = 1, 2.$$

Equilibrium

An equilibrium in this model is an acceptance policy for firms, $A^*(R, t)$, and an effort supply function for applicants, $e^*(R, k)$, that are mutual best responses to one another. We are now in a position to describe equilibrium behavior by applicants and firms under the three regulatory regimes. We require a bit more notation. Given a qualification threshold t and an acceptance capacity c, such that $F_0^{-1}(1 - c) \leq t \leq F_1^{-1}(1 - c)$, define $\hat{\pi}(t, c)$ as the solution to the equation: $F(\pi, t) = 1 - c$. That is, $\hat{\pi}(t, c)$ is the proportion of an applicant population facing the qualification threshold, t, that would need to choose high effort if the acceptance rate in that population to just equal c. Obviously,

$$\hat{\pi}(t, c) = [F_0(t) - (1 - c)]/[F_0(t) - F_1(t)].$$

Moreover, let $\Delta F(t) = [F_0(t) - F_1(t)]$, $\Delta G(k) = [G_1(k) - G_2(k)]$ and define the function: $\delta(r_2) = \omega(c - r_2)/\lambda$. Finally, let π_R^* be in fraction of group R who exert high-effort equilibrium, and let t_R^* be the qualification threshold that applies to group R in equilibrium (relevant only when firms employ a threshold policy.) It is straightforward to verify the following proposition.

Proposition 1. Equilibrium may be determined as follows.

(a) LF equilibrium is given by group-specific high effort rates (π_1^*, π_2^*) and an acceptance threshold t^* such that:

$$\hat{\pi}(t^*, c) = G(\omega\Delta F(t^*)), \quad \text{and} \quad \pi_R^* = G_R(\omega\Delta F(t^*)), \;\; R = 1, 2,$$

where firms' acceptance policy is $A^*(R, t) = 1, t \geq t^*$, and $A^*(R, t) = 0, t < t^*$, and applicants' effort supply function is $e^*(R, k) = 1$, $k \leq \omega\Delta F(t^*)$, and $e^*(R, k) = 0$, $k > \omega\Delta F(t^*)$.

(b) CS equilibrium is given by the pairs of group-specific high effort rates and acceptance thresholds, (π_1^*, t_1^*) and (π_2^*, t_2^*), where:

$$\hat{\pi}(t_R^*, r_R) = \pi_R^* = G_R(\omega\Delta F(t_R^*)), \quad R = 1, 2.$$

Firms follow the CS threshold policy $A^*(R, t) = 1, t \geq t_R^*$, and $A^*(R, t) = 0, t < t_R^*$, and applicants follow the effort supply function $e^*(R, k) = 1$, $k \leq \omega\Delta F(t_R^*)$, and $e^*(R, k) = 0$, $k > \omega\Delta F(t_R^*)$, $R = 1, 2$.

(c) CB equilibrium entails firms choosing some feasible acceptance policy and applicants choosing: $e^*(R, k) = 1$ if and only if $k \leq k^*$, else $e^*(R, k) = 0$, where k^* solves $\Delta G(k^*) = \delta(r_2)/k^*$.

Proof. The claims (a) and (b) are a transparent consequence of the definition of equilibrium. Concerning (c), as noted above, capacity and representation constraints require that applicants face the effort incentive:

$$\omega(A_1 - A_0) = \omega[c - r_2]/[\lambda(\pi_1 - \pi_2)],$$

whereas optimal behavior by applicants implies that:

$$\pi_1 - \pi_2 = \Delta G(\omega(A_1 - A_0)).$$

Combining these two equations to eliminate $(A_1 - A_0)$, and identifying k^* in (c) above with the value in equilibrium of $\omega(A_1 - A_0)$ yields the stated result. ■

The Impact of Affirmative Action on Applicant Qualifications in Equilibrium

Let us now consider how the pursuit by firms of greater representation of the disadvantaged affects the equilibrium distribution of the qualifications

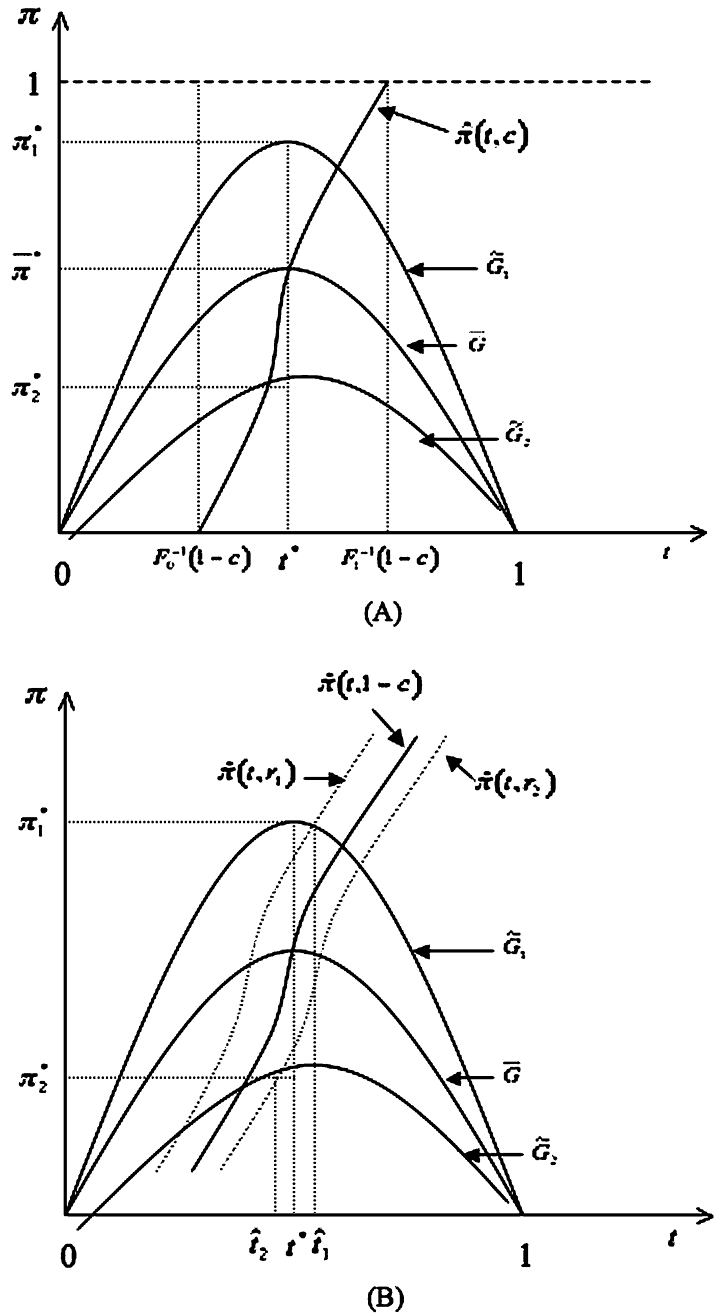

Figure 4. Laissez-faire (A) and Color-Sighted (B) Equilibrium.

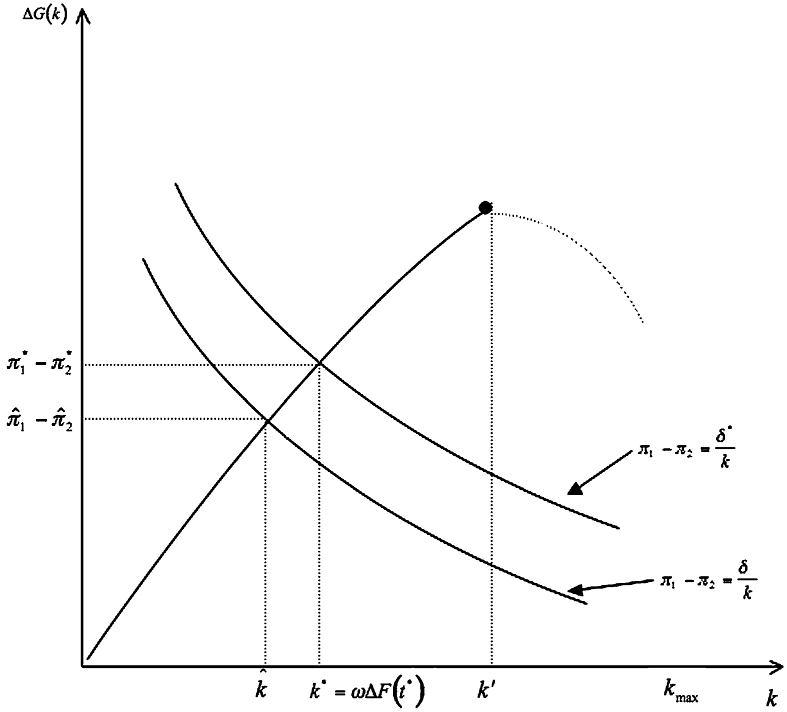

Figure 5. Color-Blind Affirmative Action Under Elitism, $\delta^* > \delta$.

presented by applicants in the two groups. We adopt the following notation: given firms' acceptance capacity, c, let $r_2^*(c)$ denote the proportion of group 2 applicants who are accepted in laissez-faire equilibrium. Thus,

$$r_2^*(c) = 1 - F(G_2[\Delta F(t^*)], t^*), \quad \text{where } t^* \text{ solves } \hat{\pi}(t^*, c) = G(\omega \Delta F(t^*)).$$

We consider the impact of weak (i.e., r_2 "close to" $r_2^*(c)$) and strong (i.e., r_2 "close to" c) affirmative action goals. We wish to distinguish two cases under which the marginally accepted applicant in the absence of any kind of affirmative action has either a low or a high level of qualifications. Specifically, let t' solve the equation: $f_1(t')/f_0(t') = 1$. If in LF equilibrium $t^* < t'$, we say that firms' equilibrium acceptance standards are loose (because the marginally accepted applicant has low qualifications), whereas if $t^* > t'$, we say firms' acceptance standards are tight (because the marginally accepted applicant has high qualifications). Examination of Figure 4 should suffice to establish the following, which we state without proof.

Proposition 2. If standards are loose in laissez-faire equilibrium, then the pursuit of sufficiently weak CS affirmative action goals leads to an increase in qualifications among the advantaged group and a decrease in qualifications among the disadvantaged group, widening the racial qualifications gap relative to laissez-faire. Moreover, tight standards in laissez-faire equilibrium imply that weak CS affirmative action decreases qualifications among the advantaged and increases qualifications among the disadvantaged, thereby narrowing the racial qualifications gap. When laissez-faire equilibrium standards are close to the margin between being loose or tight, then sufficiently strong CS affirmative action goals must decrease the qualifications of both groups.

Finally, consider the effect of color-blind affirmative action on applicant qualifications in equilibrium. To do so, we need one last definition. We will say that a selection problem is characterized by elitism if there is no acceptance policy by firms that can induce high-cost applicants to exert high effort. (Here an applicant is said to have high cost if $k > k'$, where k' solves the equation: $g_1(k')/g_2(k') = 1$.) That is, elitism obtains when the structure of the situation is such that $\omega\Delta F(t') < k'$. (Note that effort incentives are maximal at the acceptance threshold $t = t'$ and that group disparity of qualifications is greatest when the marginal applicant to choose $e^* = 1$ has cost just equal to k'.) Figure 4, together with Proposition 1(c), can be used to establish the following result:

Proposition 3. Suppose the condition of elitism obtains. Then for every affirmative action goal r_2, with $r_2^*(c) < r_2 < c$, there is a unique color-blind affirmative action equilibrium. Moreover, as r_2 rises, the level of qualifications in both groups declines, as does the qualifications gap between the groups. In the limit, as r_2 approaches c (population parity as a goal), the proportion of applicants choosing high effort in the color-blind affirmative action equilibrium approaches zero. However, there does not exist an equilibrium in this model under color-blind affirmative action that implements the representation target of population proportionality, $r_2 = c$.

In light of Figure 5, all the claims in Proposition 4 are straightforward, except the last on the impossibility of implementing an affirmative action goal of population parity in a color-blind fashion. Here is a demonstration of this result: We know that, as r_2 approaches c, applicants' incentive to exert high effort vanishes. This leads to a nonexistence result when $r_2 = c$ because, if all workers take effort $e = 0$, then the distributions of qualifications presented by workers in both groups in equilibrium would have to be identical. But then, no representation constraint would need to be imposed on firms to generate population parity. Absent such a constraint, firms would want to behave as they do under laissez-faire, accepting the fraction c of workers with the highest qualifications. Yet, were firms to do this, lower cost workers in both groups would then have an incentive to choose exert high effort, and the fact that group 2 is disadvantaged relative to group 1 would imply a failure to reach population parity in equilibrium. We conclude that when one group is disadvantaged relative to the other there can exist no color-blind equilibrium achieving population parity!

References

Akerlof, George. 1978. "The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning," 68 *American Economic Review* 8–19.

Anderson, E.J., and P. Nash. 1987. *Linear Programming in Infinite Dimensional Spaces: Theory and Applications.* New York: John Wiley and Sons.

Ayers, Ian. 1996. "Narrow Tailoring," 43 *U.C.L.A. Law Review* 1781.

Bhagwati, Jagdish. 1971. "The Generalized Theory of Distortions and Welfare," in Bhagwati, R. W. Jones, R. A. Mundell, J. Vanek, eds., *Trade, Balance of Payments and Growth.* North Holland (Kindelberger Festschrift).

Bowen, William, and Bok, Derek. 1998. *The Shape of the River: Long-Term Consequences of Considering Race in College and University Admissions.* Princeton, N.J.: Princeton University Press.

Chan, Jimmy, and Eyster, Erik. 2003. "Does Banning Affirmative Action Lower College Student Quality?," 93 *American Economic Review* 858–72.

Coate, Stephen, and Glenn C. Loury. 1993. "Will Affirmative Action Eliminate Negative Stereotypes?," 83 *American Economic Review* 1220–40.

Datcher Loury, Linda, and David Garman. 1993. "Affirmative Action in Higher Education," 83 *American Economic Review* 99–103.

Epple, Dennis, Richard Romano, and Holger Sieg. 2003, "The Practice and Proscription of Affirmative Action in Higher Education: An Equilibrium Analysis." NBER Working Paper #9799.

Fryer, Roland, and Loury, Glenn. 2003. *Categorical Redistribution in Winner-Take-All Markets.* Mimeo, Harvard University.

Holzer, Harry, and David Neumark. 2000. "Assessing Affirmative Action," 38 *Journal of Economic Literature* 483–568.

Kahlenberg, Richard D. 1996. *The Remedy: Class, Race and Affirmative Action.* New York: Basic Books.

Kane, Thomas J. 1998. "Racial & Ethnic Preferences in College Admissions," in C. Jencks and M. Phillips, eds., *The Black-White Test Score Gap.* Washington, D.C.: Brookings Institution Press.

Loury, Glenn C., Nathan Glazer, John F. Kain, Thomas Kane, Douglas Massey, Marta Tienda, Brain Bucks, et al. 2003. "Brief of Social Scientists as Amici Curiae in Support of Respondents," in *In the Supreme Court of the United States, Barbara Grutter, Petitioner, vs. Lee Bollinger et al., Respondents, No.02-241*, www.bu.edu/irsd/MichiganBrief.doc.

Sander, Richard. 1997. "Experimenting with Class-Based Affirmative Action," 47 *Journal of Legal Education* 472–503.